

M-Lab specifications for an open, independently managed broadband measurement platform

Authors:

Larry Peterson: Robert E. Kahn Professor of Computer Science at Princeton University, Director of the Princeton-hosted PlanetLab Consortium, and M-Lab Steering Committee member

Stephen Soltesz: M-Lab Platform Manager at the New America Foundation's Open Technology Institute.

Overview:

The authors, on behalf of the M-Lab Consortium, present this document in response to questions posted by the Federal Communications Commission (FCC) and others involved in the Measuring Broadband America program. These questions center on the definition of an open, transparent, independently-managed broadband measurement platform. We enumerate the requirements that should be met to achieve this, and the way in which M-Lab, as a platform designed and managed for broadband measurement by the research community, draws on years of experience to meet these requirements.

High-level requirements:

Any platform wishing to support open, verifiable, independent broadband measurement must meet these minimum criteria. The M-Lab platform meets these criteria. We review the specific ways in which M-Lab achieves this below.

1. The platform's entire source code must be freely and immediately available.
2. The platform leverages existing solutions and standards.
3. The platform has a history of including the global research community, and incorporating best practices and requirements openly proposed and verified by this community.
 - a. The platform is tuned to the needs of researchers for quality research. This is the highest bar, and meeting this requirement ensures that data collected on this platform will be suitable as the foundation of policy, consumer education, and other uses that must rely on robust results to make good choices.
4. Platform managed openly by members of the research community.
5. The platform sites do not favor any single network provider.
6. The platform is engineered at Internet scale, with consistency across deployed sites.

Requirements in detail

The following sections expand on the rationale for enforcing these requirements, and illustrate the way in which the M-Lab platform already operates to meet these criteria.

Before beginning, a quick overview to encapsulate the resources, history, and qualifications that the M-Lab platform provides and leverages:

1. M-Lab provides full-time staff with Ph.D-level experience in platform measurement who are continually monitoring, improving, and managing the operations of M-Lab's global platform.
2. M-Lab is a part of PlanetLab, and is managed in part by the Princeton-hosted PlanetLab Consortium. PlanetLab has been the state-of-the art platform for network research over the last 10 years, and this experience and concomitant improvements are all reflected in M-Lab.
3. M-Lab is a global research consortium. It leverages the many eyes and unmatched domain expertise of those at the top of their field to ensure that each piece of M-Lab works to provide a research-quality measurement environment producing robust data at Internet scale.
4. M-Lab is fully extensible. This means that anyone with the will and the resources can set up a server instance as a part of the M-Lab platform, expanding the platform and the partnership.
5. The M-Lab consortium includes organizations well-versed in operating infrastructure at Internet scale, specifically PlanetLab. These M-Lab partners have translated these years of experience into a service-level management framework for the platform and its experiments. This includes tools developed within the network operator community, like Nagios for service-level monitoring and alerting, RT for documenting and managing workflow; network elements are managed with RANCID (configuration polling) and rrdtool-based SNMP pollers (traffic monitoring). All tools used to manage the platform are open-source, leveraging improvements made to meet the requirements of the service provider community.
6. M-Lab was built with security in mind. Both the experiment hosts and the OA&M hosts are monitored for software packages that have published vulnerabilities (Fedora Security Updates for M-Lab hosts, package updates from PLC for the experiment platform, FreeBSD's "portaudit" for the OA&M hosts), and reported vulnerabilities are addressed immediately.

1. The platform's entire source code is freely and immediately available

An open-source platform serves several goals, essential to robust, scientifically verifiable measurement:

Source code availability permits third-parties to audit the platform for security, performance, and correctness. A platform built using private source systems and tools may be secure, and may perform optimally and correctly, but these properties cannot be independently confirmed or challenged by third-parties. This independent confirmation is essential when dealing with something as complex as broadband measurement.

Source code availability also permits third-parties to reproduce the platform for the sake of development, validation, or private deployments. Any platform that lacks public access to the most currently deployed source code introduces doubt regarding the foundation of its security, performance and correctness.

Measurement Lab is built on an open source platform in which the platform source code is freely and immediately available. Measurement Lab can be conceptualized as a private deployment of the PlanetLab platform, with extensive service-level for OA&M for experiments running on the platform. PlanetLab is both an open-source project and a state-of-the-art research testbed. Every piece of code necessary to build, deploy and

operate the platform is available today to anyone here:

- <https://svn.planet-lab.org>
- <https://git.planet-lab.org/>

2. The platform leverages existing solutions

Platform development and operation are extensive tasks. This cannot be overstated. Existing solutions, leveraging years of expertise in the research and academic communities, should be leveraged to meet present needs. Robust existing solutions bring with them proven history of operational success. Past performance demonstrates the consistent ability to address issues, incorporate feature requests, and continuously improve the platform.

PlanetLab, the platform underlying M-Lab, has operated continually and globally for the last decade, and now offers as an existing solution the incremental improvements and evolution gained over of ten years of experience. This includes a great breadth of experience drawing from several large scale testbeds around the world, based out of Princeton University, the Sorbonne's UPMC in Paris, and others that work together to support hundreds of experiments. Fixes and improvements from any of the existing testbeds can and do get submitted upstream to the benefit of the entire community of use and for simplified maintenance and support of each testbed.

Measurement Lab leverages the PlanetLab platform, and the experienced PlanetLab team around the globe lends its expertise and decades of operational research experience to the management and constant improvement of M-Lab. This means that M-Lab is able to deliver a platform that meets the requirements of researchers, and thus the highest standards for robust, reliable, scientifically verifiable data for all Internet stakeholders. More on PlanetLab's place as a leader in platform research can be found here: <https://www.planet-lab.org/history>.

3. The platform has a history of meeting requirements of the global research community.

The Measuring Broadband America program exists to provide accurate, comprehensive information on the performance of broadband networks in the United States. This information founds the bedrock of broadband policy and consumer understanding. Therefore, the platform chosen to conduct these measurements must have a history of meeting the highest research standards.

The PlanetLab platform has evolved to support the needs of researchers over the last decade. M-Lab is a part of the PlanetLab platform. Researchers have used the platform continuously to publish leading-edge research -- research that has pushed and expanded the global understanding of networks, and has had a tangible hand in allowing the global reach of the Internet as a crucial consumer-facing utility. A partial bibliography of papers, over two hundred, published using the PlanetLab platform is here: <https://www.planet-lab.org/biblio>. This selection includes not only research explorations of existing and upcoming network deployments and technologies, but research done to expand, improve, and stabilize the platform itself.

As a part of the PlanetLab platform, the M-Lab platform builds on technology with a decade of development, and a proven record of utility to network and distributed system researchers. The same platform is used by numerous other organizations in Japan, France, Germany, and private companies.

4. The M-Lab platform is built on years of service-provider experience

M-Lab also brings a history of operational experience with this platform, including platform management and oversight out of Princeton and the Open Technology Institute. M-Lab provides extensive service-level OA&M for the tests running on the platform, using tools and practices developed within the service provider community. This work continues to improve M-Lab's distinct model of offering more consistently provisioned and high quality capabilities such as the full-time gigabit uplink and fixed ratio of compute power to hosted experiments. The platform has grown consistently and new sites that both expand coverage and add redundancy are available all the time.

5. The platform is managed openly by members of the research community.

It is crucial for the credibility of measurement, and to avoid issues that could compromise an expensive and complex effort, that the platform be run by a neutral third party, in addition to its allowing access to source code and documentation. The M-Lab platform, as a part of PlanetLab and incorporating

Measurement Lab is operated by the New America Foundation's Open Technology Institute and Princeton University, with the active involvement of a broad consortium of partners. PlanetLab Central, operated by Princeton, works to provide in-depth support of M-Lab's operational capacity, and does so through independent administrative access to the platform. This provides a cross check monitored by a diverse network of academic partners and researchers, all of whom are involved in the M-Lab project as a means to collect and provide robust, research-quality data. Princeton has been and continues to be instrumental in diagnosing and addressing issues raised not only by SamKnows in the context of the MBA program, but across the entire platform, worldwide, regardless of hosting facility or connectivity provider.

6. Platform sites do not favor any single network provider

Networks are accurately named. They are the intermingling of many providers' infrastructure and assumptions based on open standards and shared agreements and expectations. Broadband performance is determined not by an individual provider, but by traffic's transmission across provider boundaries. It is crucial that the data that forms the foundation of policy be cognizant of this reality, and that the chosen measurement platform be operated in a way that takes this into account.

Measurement Lab has consistently managed, consistently deployed sites throughout the world, and the network providers at these sites are also varied. While the FCC's interest in Me- Lab focuses on connectivity within the USA, the fact the Measurement Lab hosts sites outside the US is an indication of the Platform's lack of bias toward any one country or ISP.

7. Engineered at Internet scale

The Internet is a vast and dynamic system. And, as a result, the Internet embodies considerable complexity. Any platform designed to support continuous measurement needs to be able to quickly adapt and respond while still maintaining consistency of services.

M-Lab operates at this scale, backed by the large and vibrant community around the open source PlanetLab stack and the deep history of that project's response to new and constantly evolving research needs coupled with the consistently monitored and managed resources that are as uniformly provisioned as is possible for an effort of this size that has its own not inconsiderable history over many years.

Expanding the M-Lab platform to improve redundancy

The Commission and others involved in the MBA program have expressed a desire to expand the measurement platform to improve redundancy, toward a goal of uninterrupted data-collection. The M-Lab Consortium agrees with this goal. Currently, some aspects of it are met in that each site is deployed with three servers for site-level redundancy, and experiment traffic can be diverted to alternate sites without incident for latency-insensitive experiments. M-Lab is committed to continuous improvement of its operational practices, and is more than willing to work with the FCC toward shared outcomes. However, the infrastructure collecting data for the program should meet the requirements enumerated throughout this document. M-Lab meets these requirements, and is specifically designed to be easily extensible. This means that anyone with the resources can provide M-Lab servers that meet the requirements and that are fully managed by the M-Lab Consortium.

The M-Lab consortium, and the broader research and academic community strongly encourage the Commission and those within the program to meet the goal of platform redundancy by donating servers to the existing M-Lab platform.

The specifications and requirements for donation are available here:

<http://measurementlab.net/sites/default/files/HowtoContributeToM-LabServerInfrastructure.pdf>

To discuss possible donations, please contact Thomas Gideon (gideon@newamerica.net) and Stephen Soltesz (soltesz@newamerica.net).